

What is claimed is:

1. In a data collection agent, a method for processing storage information concerning

5 storage entities in a storage area network environment, the method comprising:

receiving a data collection policy identifying storage information to be collected concerning a set of storage entities in the storage area network environment;

collecting storage information concerning the set of storage entities according to the data collection policy;

10 processing the collected storage information to produce at least one summary record of the storage information for the set of storage entities, the at least one summary record containing a summary of the collected storage information for the set of storage entities calculated according to the data collection policy; and

transferring the at least one summary record from the data collection agent to a
15 storage management application for presentation to a user of the storage management application.

2. The method of claim 1 wherein receiving a data collection policy comprises:

receiving a collection level identifying a set of files corresponding to the set of
20 storage entities from which storage information is to be collected; and

receiving a collection period value indicating how often to perform collection of the storage information concerning the set of storage entities.

3. The method of claim 2 wherein receiving a collection level corresponding to the set of
25 storage entities from which storage information is to be collected comprises:

receiving at least one of a host collection level, a file system collection level and a file set collection level; and

based on the collection level, identifying, as the set of storage entities, a set of file systems containing at least one file from which storage information is to be collected.

30

4. The method of claim 2 wherein receiving a collection level comprises:

receiving a selection of a host collection level identifying at least one host computer system that maintains access to a set of file systems from which storage information is to be collected; and

5 identifying, as the set of storage entities from which storage information is to be collected, files within a set of file systems accessible within data storage systems associated with the at least one host computer system identified by the host collection level.

5. The method of claim 2 wherein receiving a collection level comprises:

10 receiving a selection of a file system collection level identifying at least one file system that maintains access to the set of files from which storage information is to be collected; and

15 identifying, as the set of storage entities from which storage information is to be collected, files within the identified at least one file system corresponding to the file system collection level.

6. The method of claim 2 wherein receiving a collection level comprises:

20 receiving a selection of a file set collection level identifying at least one specific file from which storage information is to be collected, the selected at least one specific file indicating the set of storage entities from which storage information is to be collected.

7. The method of claim 6 wherein receiving a selection of a file set collection level comprises receiving file filter criteria including at least one of:

- 25
- i) at least a portion of a name of a file;
 - ii) an owner of a file;
 - iii) an age of a file; and
 - iv) a size of a file,

30 the file filter criteria identifying file characteristics of files to be included within the set of storage entities from which storage information is to be collected.

8. The method of claim 3 wherein collecting storage information concerning the storage entity according to the data collection policy comprises:

detecting occurrence of the collection period;

in response to the occurrence of the collection period, parsing the set of file

5 systems to query each of the at least one files within directories in the set of file systems identified based on the collection level to obtain file metadata associated with each file; and

based on the file metadata for each file detected in each of the set of file systems during parsing, collecting the storage information related to ages and sizes of files and
10 directories for each storage entity associated with at least one of the host collection level, the file system collection level and a file set collection level.

9. The method of claim 8 wherein collecting the storage information related to ages and sizes of files and directories for each storage entity comprises calculating at least one of:

15 i) a running total number of files detected in each directory in the file system;

ii) a running total of the amount of data maintained within the each file detected;

iii) a running total of the age of each file;

20 iv) for each directory detected in each of the set of file systems during parsing, a running total of the size of all files in that directory;

v) a number N of the largest files for the collection level; and

vi) a number M of the oldest files for the collection level.

25 10. The method of claim 9 wherein processing the collected storage information to produce a summary record of the storage information for the set of storage entities comprises:

calculating age characteristics related to files and directories within the set of file systems identified based on collected storage information for the collection level; and

30 calculating size characteristics related to files and directories within the set of file systems identified based on collected storage information for the collection level.

11. The method of claim 10 wherein the collection level is a host collection level; and wherein calculating age averages related files and directories within the set of file systems identified based on collected storage information for the host collection level
5 comprises:

for all file systems accessible to host computer systems associated with the host collection level, producing an age summary record indicating an age of files within the file systems for each host for a plurality of age ranges and an amount of storage the files within each age range occupy and a list of the top M oldest files; and

10 wherein calculating size averages related files and directories within the set of file systems identified based on collected storage information for the host collection level comprises:

for all file systems accessible to host computer systems associated with the host collection level, producing a size summary record indicating an total amount of files
15 stored within the file systems for each host, at least N of the largest files accessible to the host.

12. The method of claim 9 wherein the collection level is a file system collection level; and wherein calculating age averages related files and directories within the set of file
20 systems identified based on collected storage information for the file system collection level comprises:

for all file systems associated with the file system collection level, producing an age summary record indicating an age of files within the file systems for a plurality of age ranges and an amount of storage the files within each age range occupy and a list of
25 the top M oldest files with the file systems; and

wherein calculating size averages related files and directories within the set of file systems identified based on collected storage information for the file system collection level comprises:

for all file systems associated with the file system collection level, producing a
30 size summary record indicating an total amount of files stored within the file systems, at least N of the largest files accessible within the file systems.

13. The method of claim 1 wherein the data collection policy received defines a host collection level and wherein collecting storage information concerning the set of storage entities according to the data collection policy comprises:

5 for each host defined by the host collection level, collecting storage information related to the sizes and ages of files and directories within file systems accessible to that host;

 and wherein processing the collected storage information to produce at least one summary record of the storage information for the set of storage entities comprises:

10 for each host defined by the host collection level, producing a size summary record and an age summary record by applying averaging functions to the storage information related to the sizes and ages of files and directories within file systems accessible to that host.

15 14. The method of claim 1 wherein the data collection policy received defines a file system collection level and wherein collecting storage information concerning the set of storage entities according to the data collection policy comprises:

 for each file system defined by the file system collection level, collecting storage information related to the sizes and ages of files and directories within that file system;

20 and wherein processing the collected storage information to produce at least one summary record of the storage information for the set of storage entities comprises:

 for each file system defined by the file system collection level, producing a size summary record and an age summary record by applying averaging functions to the storage information related to the sizes and ages of files and directories within that file system.

25 15. The method of claim 1 wherein the data collection policy received defines a file set collection level indicating a specific set of files for which storage information is to be collection and wherein collecting storage information concerning the set of storage entities according to the data collection policy comprises:

for each file defined by the file set collection level, collecting storage information related to the size and age of that file;

and wherein processing the collected storage information to produce at least one summary record of the storage information for the set of storage entities comprises:

5 for all files defined by the file set collection level, producing a size summary record and an age summary record by applying averaging functions to the storage information related to the sizes and ages of the files identified by the file set collection level.

10 16. A computerized device comprising:

 a network interface coupled to a network;

 a storage interface coupling a data storage system that maintain storage entities to the computerized device;

 a processor;

15 a memory system;

 an interconnection mechanism coupling the network interface, the storage interface, the processor and the memory system;

 wherein the memory system is encoded with an agent application that, when executed on the processor, produces an agent process that processes storage information
20 concerning storage entities in a storage area network environment that are associated with the computerized device by causing the computerized device to perform the operations of:

 receiving, via the network interface, a data collection policy identifying storage information to be collected concerning a set of storage entities that the computerized
25 device maintains in the data storage system;

 collecting storage information concerning the set of storage entities according to the data collection policy;

 processing the collected storage information to produce at least one summary record of the storage information for the set of storage entities, the at least one summary
30 record containing a summary of the collected storage information for the set of storage entities calculated according to the data collection policy; and

transferring, via the network interface, the at least one summary record from the data collection agent to a storage management application for presentation to a user of the storage management application.

- 5 17. The computerized device of claim 16 wherein when the agent process causes the computerized device to perform the operation of receiving a data collection policy, the agent process causes the computerized device to perform the operations of:

receiving a collection level identifying a set of files corresponding to the set of storage entities from which storage information is to be collected; and

- 10 receiving a collection period value indicating how often to perform collection of the storage information concerning the set of storage entities.

18. The computerized device of claim 17 wherein when the agent process causes the computerized device to perform the operation of receiving a collection level

- 15 corresponding to the set of storage entities from which storage information is to be collected, the agent process causes the computerized device to perform the operations of:

receiving at least one of a host collection level, a file system collection level and a file set collection level; and

- 20 based on the collection level, identifying, as the set of storage entities, a set of file systems containing at least one file from which storage information is to be collected.

19. The computerized device of claim 17 wherein when the agent process causes the computerized device to perform the operation of receiving a collection level, the agent process causes the computerized device to perform the operations of:

- 25 receiving a selection of a host collection level identifying at least one host computer system that maintains access to a set of file systems from which storage information is to be collected; and

- identifying, as the set of storage entities from which storage information is to be collected, files within a set of file systems accessible within data storage systems
30 associated with the at least one host computer system identified by the host collection level.

20. The computerized device of claim 17 wherein when the agent process causes the computerized device to perform the operation of receiving a collection level, the agent process causes the computerized device to perform the operations of:

5 receiving a selection of a file system collection level identifying at least one file system that maintains access to the set of files from which storage information is to be collected; and

 identifying, as the set of storage entities from which storage information is to be collected, files within the identified at least one file system corresponding to the file
10 system collection level.

21. The computerized device of claim 17 wherein when the agent process causes the computerized device to perform the operation of receiving a collection level, the agent process causes the computerized device to perform the operation of:

15 receiving a selection of a file set collection level identifying at least one specific file from which storage information is to be collected, the selected at least one specific file indicating the set of storage entities from which storage information is to be collected.

20 22. The computerized device of claim 21 wherein when the agent process causes the computerized device to perform the operation of receiving a selection of a file set collection level, the agent process causes the computerized device to perform the operation of receiving file filter criteria including at least one of:

- i) at least a portion of a name of a file;
25 ii) an owner of a file;
 iii) an age of a file;
 iv) a size of a file; and

 wherein the file filter criteria identifies file characteristics of files to be included within the set of storage entities from which storage information is to be collected.

23. The computerized device of claim 19 wherein when the agent process causes the computerized device to perform the operation of collecting storage information concerning the storage entity according to the data collection policy, the agent process causes the computerized device to perform the operation of:

- 5 detecting occurrence of the collection period;
 in response to the occurrence of the collection period, parsing the set of file systems to query each of the at least one files within directories in the set of file systems identified based on the collection level to obtain file metadata associated with each file; and
10 based on the file metadata for each file detected in each of the set of file systems during parsing, collecting the storage information related to ages and sizes of files and directories for each storage entity associated with at least one of the host collection level, the file system collection level and a file set collection level.

- 15 24. The computerized device of claim 23 wherein when the agent process causes the computerized device to perform the operation of collecting the storage information related to ages and sizes of files and directories for each storage entity, the agent process causes the computerized device to perform the operation of calculating at least one of:

- 20 i) a running total number of files detected in each directory in the file system;
 ii) a running total of the amount of data maintained within the each file detected;
 iii) a running total of the age of each file;
 iv) for each directory detected in each of the set of file systems during
25 parsing, a running total of the size of all files in that directory;
 v) a number N of the largest files for the collection level; and
 vi) a number M of the oldest files for the collection level.

- 30 25. The computerized device of claim 24 wherein when the agent process causes the computerized device to perform the operation of processing the collected storage

information to produce a summary record of the storage information for the set of storage entities, the agent process causes the computerized device to perform the operations of:

- calculating age characteristics related to files and directories within the set of file systems identified based on collected storage information for the collection level; and
- 5 calculating size characteristics related to files and directories within the set of file systems identified based on collected storage information for the collection level.

26. The computerized device of claim 25 wherein the collection level is a host collection level; and wherein when the agent process causes the computerized device to perform the
10 operation of calculating age averages related files and directories within the set of file systems identified based on collected storage information for the host collection level, the agent process causes the computerized device to perform the operations of:

- for all file systems accessible to host computer systems associated with the host collection level, producing an age summary record indicating an age of files within the
15 file systems for each host for a plurality of age ranges and an amount of storage the files within each age range occupy and a list of the top M oldest files; and

wherein when the agent process causes the computerized device to perform the operation of calculating size averages related files and directories within the set of file systems identified based on collected storage information for the host collection level, the
20 agent process causes the computerized device to perform the operations of:

- for all file systems accessible to host computer systems associated with the host collection level, producing a size summary record indicating an total amount of files stored within the file systems for each host, at least N of the largest files accessible to the
25 host.

27. The computerized device of claim 24 wherein the collection level is a file system collection level; and wherein when the agent process causes the computerized device to perform the operation of calculating age averages related files and directories within the set of file systems identified based on collected storage information for the file system
30 collection level, the agent process causes the computerized device to perform the operations of:

for all file systems associated with the file system collection level, producing an age summary record indicating an age of files within the file systems for a plurality of age ranges and an amount of storage the files within each age range occupy and a list of the top M oldest files with the file systems; and

5 wherein when the agent process causes the computerized device to perform the operation of calculating size averages related files and directories within the set of file systems identified based on collected storage information for the file system collection level, the agent process causes the computerized device to perform the operations of:

10 for all file systems associated with the file system collection level, producing a size summary record indicating an total amount of files stored within the file systems, at least N of the largest files accessible within the file systems.

28. The computerized device of claim 16 wherein the data collection policy received defines a host collection level and wherein when the agent process causes the
15 computerized device to perform the operation of collecting storage information concerning the set of storage entities according to the data collection policy, the agent process causes the computerized device to perform the operations of:

20 for each host defined by the host collection level, collecting storage information related to the sizes and ages of files and directories within file systems accessible to that host;

 and wherein when the agent process causes the computerized device to perform the operation of processing the collected storage information to produce at least one summary record of the storage information for the set of storage entities, the agent process causes the computerized device to perform the operations of:

25 for each host defined by the host collection level, producing a size summary record and an age summary record by applying averaging functions to the storage information related to the sizes and ages of files and directories within file systems accessible to that host.

30 29. The computerized device of claim 16 wherein the data collection policy received defines a file system collection level and wherein when the agent process causes the

computerized device to perform the operation of collecting storage information concerning the set of storage entities according to the data collection policy, the agent process causes the computerized device to perform the operations of:

- 5 for each file system defined by the file system collection level, collecting storage information related to the sizes and ages of files and directories within that file system;
- and wherein when the agent process causes the computerized device to perform the operation of the collected storage information to produce at least one summary record of the storage information for the set of storage entities, the agent process causes the computerized device to perform the operations of:
- 10 for each file system defined by the file system collection level, producing a size summary record and an age summary record by applying averaging functions to the storage information related to the sizes and ages of files and directories within that file system.

- 15 30. The computerized device of claim 16 wherein the data collection policy received defines a file set collection level indicating a specific set of files for which storage information is to be collection and wherein when the agent process causes the computerized device to perform the operation of collecting storage information concerning the set of storage entities according to the data collection policy, the agent
- 20 process causes the computerized device to perform the operations of:

- for each file defined by the file set collection level, collecting storage information related to the size and age of that file;
- and wherein when the agent process causes the computerized device to perform the operation of processing the collected storage information to produce at least one
- 25 summary record of the storage information for the set of storage entities, the agent process causes the computerized device to perform the operations of:
- for all files defined by the file set collection level, producing a size summary record and an age summary record by applying averaging functions to the storage information related to the sizes and ages of the files identified by the file set collection
- 30 level.

31. A computer program product having a computer readable medium operable to store computer program logic embodied in computer program code encoded thereon for processing timer events in a services infrastructure, the computer program code, when executed on a processor in a computerized device, causing the computerized device to perform the operations of:

receiving a data collection policy identifying storage information to be collected concerning a set of storage entities in the storage area network environment;

collecting storage information concerning the set of storage entities according to the data collection policy;

processing the collected storage information to produce at least one summary record of the storage information for the set of storage entities, the at least one summary record containing a summary of the collected storage information for the set of storage entities calculated according to the data collection policy; and

transferring the at least one summary record from the data collection agent to a storage management application for presentation to a user of the storage management application.